CLAIMS

What is claimed is:

1. A method of selecting a physical memory locality in a multiprocessor system, the method comprising:
providing a data structure including sets of equidistant localities; and
determining a preferred locality using a pointer to a locality within said data structure.

2. The method of claim 1, further comprising:
receiving an initial locality request including an indication of a search policy; and
forming the data structure using physical memory localities within the system and using the search policy.

3. The method of claim 2, wherein the physical memory localities include local memories at cells in the system.

4. The method of claim 3, wherein the search policy comprises a "closest first" policy.

5. The method of claim 3, wherein the physical memory localities further includes interleaved memory in the system.

6. The method of claim 5, wherein the search policy comprises an "interleaved first" type of policy.

7. The method of claim 1, wherein the determination of the preferred locality is is performed using a get "best"/"next best" iteration procedure.

8. The method of claim 1, wherein the pointer is rotated amongst localities within a current equidistant set so as to provide for round-robin type selection amongst those equidistant localities.

9. The method of claim 1, wherein the determination of the preferred locality includes changing to a next equidistant set if there is no memory available in any locality of a current equidistant set.

10. The method of claim 9, further comprising returning an indication that no locality is available if no locality within any of the equidistant sets has sufficient memory.

11. A multiprocessor computing system, the system comprising:
multiple symmetric multiprocessing (SMP) nodes;
multiple central processing units (CPUs) at each SMP node;
a memory control unit at each SMP node which is coupled to each CPU at that SMP node;
shared memory at each SMP node which is accessible by way of the memory control unit at that SMP node;
a switching system coupled to the memory control units so as to interconnect the multiple SMP nodes;
an operating system running on the CPUs;
a virtual memory (VM) fault handler within the operating system; and
a VM locality module within the operating system; and
a data structure of physical memory localities,
wherein the VM locality module determines a preferred locality using a pointer to a locality within the data structure.

12. The system of claim 11, wherein the data structure comprises sets of equidistant localities.

13. The system of claim 12, wherein the preferred locality is determined using a "closest first" search policy.

14. The system of claim 13, wherein the data structure comprises a first set including a closest local memory locality and one or more other sets of equidistant localities.

15. The system of claim 14, wherein the other sets include an interleaved memory locality.

16. The system of claim 11, wherein the shared memory includes both local memory and interleaved memory.

17. The system of claim 16, wherein the preferred locality is determined using an "interleaved first" search policy.

18. The system of claim 17, wherein the data structure comprises a first set including an interleaved memory locality and a set including local memory localities.

19. A data structure for use in selecting a physical memory locality in a multiprocessor system, the data structure being configured in accordance with a search policy and comprising multiple sets of equidistant physical memory localities under the search policy.

20. The data structure of claim 19, wherein the search policy comprises an "interleaved first" policy, and wherein a first set comprises an interleaved memory locality.

21. The data structure of claim 19, wherein the search policy comprises an "closest first" policy, and wherein a first set comprises a most rapidly accessible memory locality.

22. An operating system for a multiprocessor computing system, the operating system comprising:

a virtual memory manager configured for extending a memory space beyond limits of a physical address space;

a virtual memory locality module configured to rapidly select a physical memory locality in the system; and

a virtual memory fault handler configured to interrupt execution of the virtual memory manager when a page fault occurs and to utilize the virtual memory locality module to determine the physical memory locality from which to allocate memory in response to the page fault.